

REMARKS

The present response is intended to be fully responsive to all points raised by the Examiner in the Advisory Action and is believed to place the application in condition for allowance. Favorable reconsideration and allowance of the application is respectfully requested.

The application as examined includes claims 27 – 38, 40 – 51, and 71 – 78. In the present response, claims 27, 42, and 49 are amended; no additional claims are canceled, and no new claims are added.

Support for amendments to the claims is set forth hereinbelow, with reference to the application as published in U.S. Patent Application Publication No. 2003/0130831:

The claimed feature recited in amended independent claims 27, 42, and 49 as “a combination identifier which includes a string of literals identifying a particular outcome of said test program” is disclosed in paragraph [0227].

Claims 27 – 38 and 40 – 51 stand rejected under 35 U.S.C. §103(a) as being unpatentable over *Genesys-MP: User's Guide* (“Genie”) in view of “A Simulation Based Approach to Architectural Verification of Multiprocessor Systems” (“Saha”).

The present invention provides a computer program product for validating a multi-processor design by simulating program execution for a test program having at least two simulated processes which access mutually-dependent non-adjacent resources, and by creating one or more tagged value-lists incorporating a set of non-unique values associated with a combination identifier which includes a string of literals identifying a particular outcome of the test program. The contents of the resources are then compared with the non-unique values in the tagged value-lists to validate the processor design for the test program.

Genie describes the use and operation of the “Genesys” multi-processor test system.

Saha describes verifying coherency in weakly ordered shared memory multiprocessor systems through the use of “data coloring”.

In the Advisory Action, it is stated that:

‘The prior art teaches that “valid or possible set of values are maintained for each referenced location on a per processor basis.”

This is equivalent to the claimed “non-unique values.” Saha’s teaching of “uniquely colored data” or unique data for each store is immaterial, based on the aforementioned interpretation of “non-unique values”.’

The Applicant agrees that the prior art of Saha teaches that “valid or possible set of values are maintained for each referenced location on a per processor basis.”

However, the Applicant respectfully disagrees with the Examiner’s interpretation that this is equivalent to the claimed non-unique values. The Applicant also respectfully disagrees that Saha’s teaching of “uniquely colored data” or unique data for each store is “immortal”. This is detailed as follows:

First, Saha clearly teaches that “*valid sets* and the use of uniquely colored data for stores” are separate and distinct from one another (Saha, page 35, Section 2. third paragraph). Saha Section 2.2 “Data Coloring and Valid Sets” (Saha, page 36) stipulates “**two** fundamental concepts”:

“Data is unique for each store. If all the different stores write differently colored data, then a read back of a certain data value can be used to unambiguously determine the completions of a store in the memory hierarchy. Load completion messages are always exact at the processor interface since the data must be loaded into some register.”

and

“Valid or possible set of values are maintained for each referenced location on a per processor basis. This reference can be made by any processor and via any instruction.”

That is, the uniqueness of data for stores is guaranteed by data coloring, which is entirely separate from maintaining valid or possible set of values on a per processor basis.

Data coloring is used specifically because it uniquely identifies the store operation which places data in a particular location, regardless of which processor executes the store operation. Absolute uniqueness of the data is the foundation of data coloring. If different processors were allowed to write non-unique data into a location, this would defeat the purpose of data coloring.

Saha explains this in Section 2, third paragraph (Saha page 35):

“In weakly consistent systems, the exact values of shared memory locations cannot be computed statically because of the data races that can result between competing access to the same location from different processors. As a result, other prevalent architectural verification approaches ... mostly permit false sharing. ... The method presented in this paper allows true sharing access from different processors and does not require dynamic monitoring of transactions or exact completion times of stores. The dynamic transaction and coherence protocol monitoring is avoided by introducing the concept of *valid sets* and the use of uniquely colored data for stores. Since the verification algorithm is based on observing only the results of execution of memory read *instructions*, exact store completion times are not required. ...”

In other words, Saha handles the race conditions between different processors in weakly consistent systems by *uniquely* distinguishing store operations with *uniquely colored data for stores*. Thus, Saha can ignore the dynamic transaction process that stores the data (and

can ignore timing issues), because the stored data is “colored” by adding unique information to the stored data that uniquely identifies which operation stored the data. To run the verification, Saha merely has to read the memory locations — the uniquely-colored data reveals precisely which store operation stored the data.

Once again, the uniqueness of data coloring is absolutely essential. If different processors could store non-unique values at the same location, then Saha would not be able to resolve race conditions between the processors, which is Saha’s primary goal.

Saha’s teaching of “uniquely colored data” or unique data for each store is material. The fact that “valid or possible set of values are maintained for each referenced location on a per processor basis” does not negate or override the uniqueness of Saha’s data coloring.

The Applicant therefore respectfully maintains that the arguments put forth in the response filed July 27, 2010, are valid, and that the cited prior art of Genie and Saha, both individually and in combination, fail to disclose or reasonably suggest the claims of the present application. In addition to the *non-unique* values featured in the present invention versus the *unique* values featured in Saha, the tagged value lists of the present invention are structurally-distinct from Saha’s data coloring.

While continuing to respectfully traverse the 35 U.S.C. §103(a) rejection, the Applicant is amending the independent claims of the present invention to recite “a combination identifier which includes a string of literals identifying a particular outcome of said test program”. The Applicant respectfully submits that the art of record fails to disclose or reasonably suggest at least a tagged value list, wherein the tag is a combination identifier which includes a string of literals identifying a particular outcome of said test program.

The Applicant therefore respectfully submits that amended independent claims 27, 42, and 49 are patentable over the art of record. The remainder of the claims each depend directly or ultimately from the amended independent claims, and therefore are also allowable over the art of record.

In view of the foregoing remarks, all of the claims are believed to be in condition for allowance. Favorable reconsideration and allowance of the application is respectfully requested.

Please charge any fees associated with this response to Deposit Account
09-0468.

Respectfully submitted,

By: /Suzanne Erez/
Suzanne Erez
Reg. No. 46,688
Phone No. 1-888-260-5928

Date: 12 September 2010
IBM Corporation
Intellectual Property Law Dept.
P. O. Box 218
Yorktown Heights, New York 10598